(12) UK Patent Application (19) GB (11) 2 305 811 (13) A

(43) Date of A Publication 16.04.1997

(21) Application No 9519613.5

(22) Date of Filing 26.09.1995

(71) Applicant(s)
Northern Telecom Limited

(Incorporated in Canada - Quebec)

World Trade Center of Montreal,
380 St Antoine Street West,8th Floor, Montreal,
Quebec H2Y 3Y4, Canada

(72) Inventor(s)
Kevin Lewis
Bertrand Gerard Christian Mace

(74) Agent and/or Address for Service
Nortel Patents
London Road, HARLOW, Essex, CM17 9NA,
United Kingdom

(51) INT CL⁶
H04M 3/36 , H04Q 3/00

(52) UK CL (Edition O)
H4K KFM

(56) Documents Cited
GB 2253970 A   EP 0583135 A2   US 5216591 A
US 5067148 A

(58) Field of Search
UK CL (Edition N) H4K KFM
INT CL⁶ H04M , H04Q
ONLINE: WPI; INSPEC.

(54) Traffic routing in a telecommunications network

(57)   In a telecommunications network comprising a plurality of nodes (N1,N2...)interconnected by communications links (L1,L2...) each incorporating one or more communications channels, traffic routing between network nodes is determined by providing a map or model of the network and determining in the model first and second nodes between which a communications path or route is to be established across the network. A set or colony of exploratory agents is created at the first node, and explores the model network via the links and nodes between the first and second nodes, each agent recording the results of its exploration. From this exploration a route favoured by a significant number of the agents is determined and is set up on the telecommunications network.
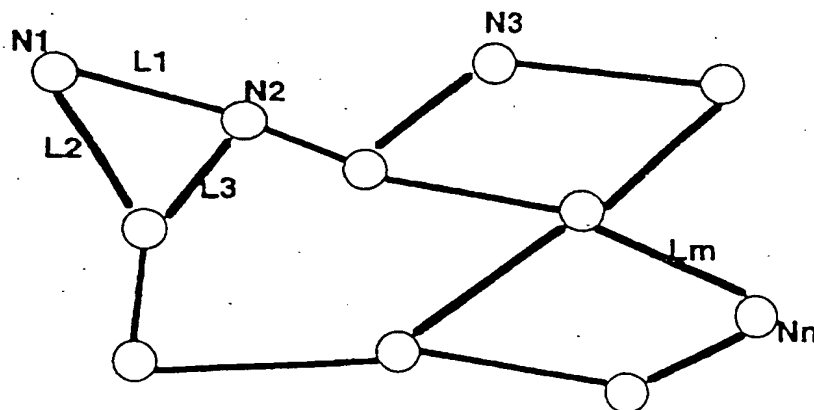
Fig.1

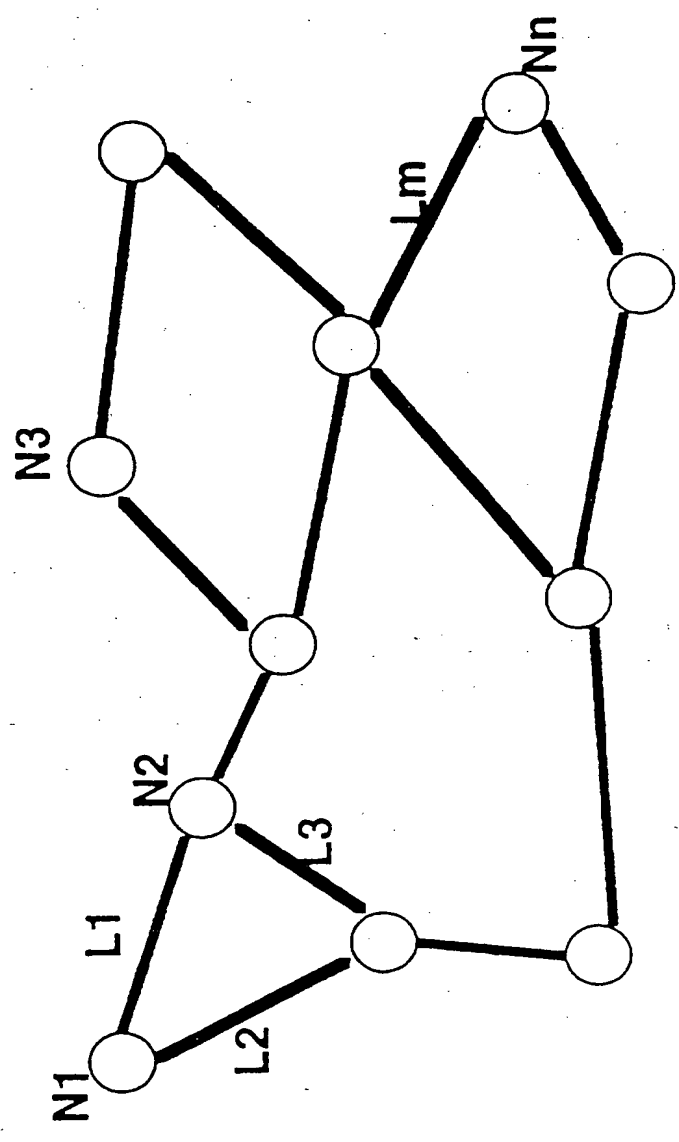GB 2 305 811 A

Fig.1
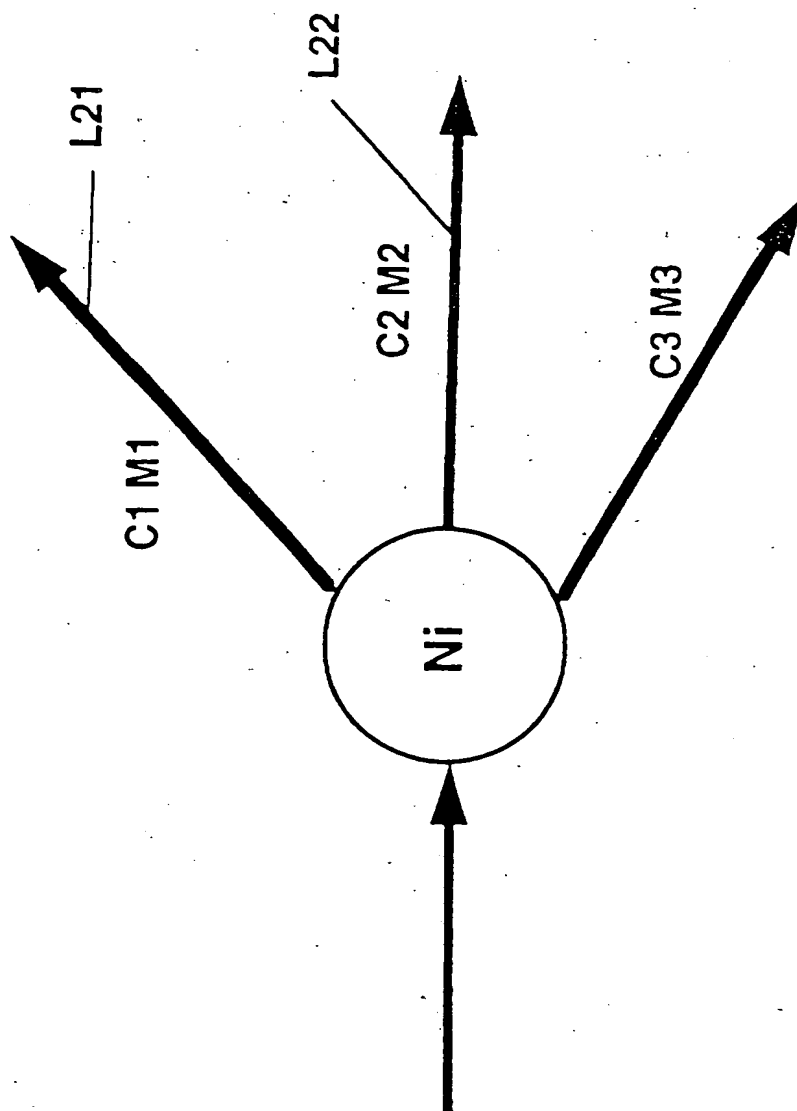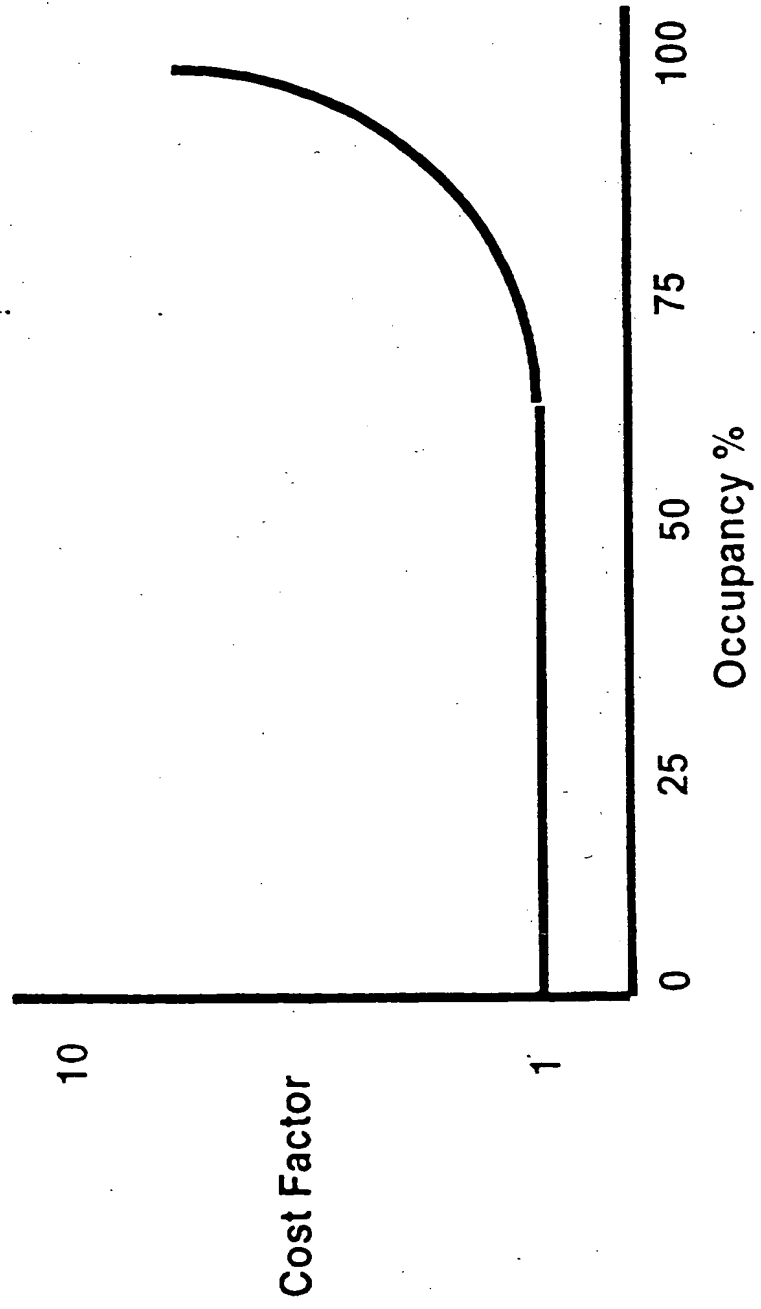
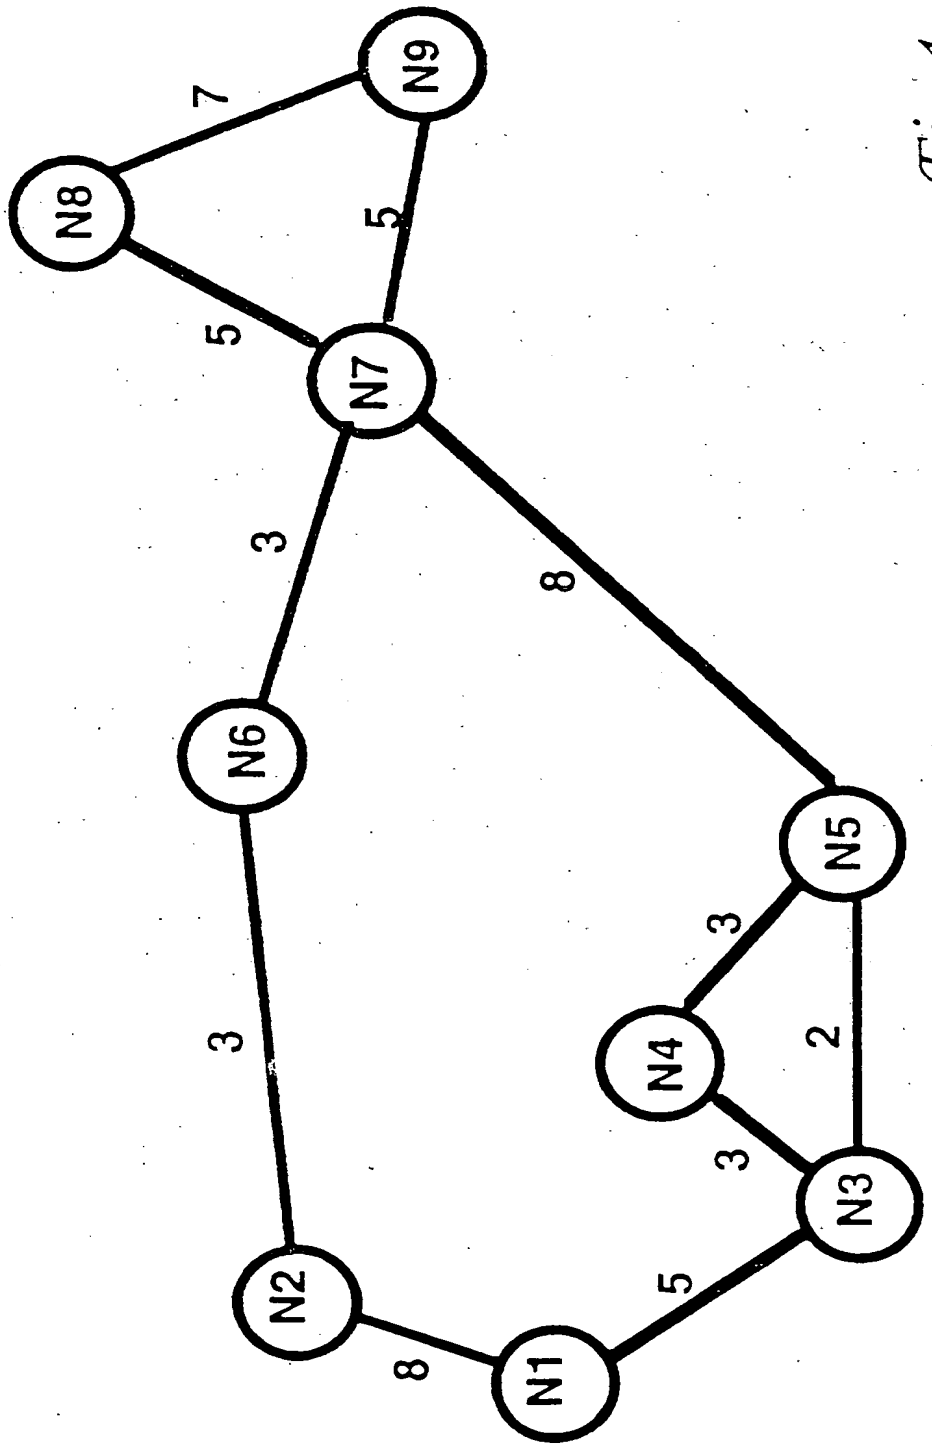C1 M1   L21

C2 M2   L22

C3 M3

Ni

Fig.2

Fig.3

Fig.4

TRAFFIC ROUTING IN A TELECOMMUNICATIONS NETWORK

This invention relates to the routing of telecommunications traffic in a network comprising a plurality of nodes interconnected by paths or links. In particular, the invention relates to a method and apparatus for operating a network by defining a traffic routing plan for such a network and to a

5   method of routing traffic via the routing plan.

As telecommunications networks become more complex and the traffic load increases there is a need to determine the routing of traffic within a network so as to provide efficient utilisation of communication channels

10   used, and to minimise the total cost to the system operator. A particular problem that frequently confronts network operators is a business customer request to provide a number of permanent virtual circuits between two network nodes. To provide these circuits the operator must choose a route across the network which takes account of cost and of

15   available bandwidth. At present, traffic routing is determined largely on an empirical basis as a full analysis of a complex network to obtain an analytical solution is an intractable problem. It is thus very difficult for an operator to ensure that a network is used with a high degree of efficiency under all traffic conditions.

20

An object of the present invention is to minimise or to overcome this disadvantage.

A further object of the invention is to provide an improved method of

25   determining routing information in a telecommunications network.

According to one apect of the invention there is provided a method of operating a telecommunications network whereby to determine traffic routing, said network comprising a plurality of nodes interconnected by

30   communications links and each said link incorporating one or more communications channels, wherein the method includes, providing a map or model of the network, determining in the model those nodes between which a communications path or route is to be established across the

network, creating a set or colony of exploratory agents at a said node, permitting the agents at said node to explore the model network via the links and nodes, recording the results of said exploration, and determining from said results a route favoured by a significant number of said agents, 5 and setting up said route corresponding to that path on the telecommunications network.

According to another aspect of the invention there is provided apparatus for determining traffic routing in a telecommunications network comprising 10 a plurality of nodes interconnected by communications links, each said link incorporating one or more communications channels, the apparatus including, means for providing a map or model of the network, means for determining in the model first and second nodes between which a communications path or route is to be established across the network, 15 means for creating a set or colony of exploratory agents at said first node, said agents being adapted to explore the model network via the links and nodes between said first and second nodes, means for recording the results of said exploration, and means for determining from said results a route favoured by a significant number of said agents whereby to set up 20 said route on the telecommunications network.

The exploratory agents are individually of very limited intelligence. Initially they explore the network selecting links based on the link cost, but as the search develops paths are found which are traversed preferentially by the 25 agents. The process may be employed by a network operator to allocate circuits in response to customer requests.

The method is of particular application to ATM networks, but it is not of course limited to networks of this type.

30

An embodiment of the invention will now be described with reference to the accompanying drawings in which:

Figure 1 is a schematic diagram of a telecommunications network 35 comprising a plurality of nodes interconnected by links;

Figure 2 shows further detail of a node of the network of figure 1;

Figure 3 illustrates a typical relationship between link occupancy and cost in the network of figure 1; and

Figure 4 shows a simple modle network used for demonstration purposes.

Referring to Figure 1, the network comprises a number of nodes N1, N2 ....Nn interconnected by communications link L1, L2 .....Lm. Each link has a respective capacity of a number of channels and also has a respective operating cost. Traffic between two network nodes is carried over a path comprising a set of links between those nodes and, in general, there will be a large number of possible paths between those nodes. For each path, the cost of that path is the total of the costs of each link comprising the path. It will be appreciated that for exploratory purposes, figure 1 shows a much simplified representation of the network which will normally comprise a large number of nodes and links.

The choice of a particular route between a pair of nodes will, in general, depend on considerations of the bandwidth required to accommodate the traffic and the cost to the network operator. However, when determining a route, i.e. selecting a series of links between selected nodes, existing traffic must also be considered to prevent the situation where some links have reached their full capacity whilst others are little used. This is generally referred to as load balancing.

A traffic routing path in the network of Figure 1 is determined by the following process. A mode or map, i.e. a software model, of the network is prepared, each link of the model being allocated a bandwidth factor and a cost factor C determined from the real network. The end nodes of the derived path are identified and the bandwidth or the number of communications channels required to carry the traffic on the path are determined. In some applications the model may be simplified at this stage by deleting any links whose bandwidth is insufficient to carry the required traffic.

The next stage is the insertion of a set or colony of software exploratory agents at one of the end nodes of the path. These agents are simple devices having very little intelligence and which explore the network passing from node to node via the links therebetween until the destination

node is reached. On this initial exploratory journey, the agents have the following characteristics.

1. Each time a link is traversed and each time a node is visited, the agent adds the identity of that link or node to a tabu list. This prevents a link being traversed or a node visited more than once and thus prevents closed loops or cycles within the path. The tabu list comprises an ordered search of the links. If an agent is unable to select a next link, then it backtracks to a node where further choices are possible.

2. When an agent arrives at a node it selects a link from those links leaving the node on an essentially random basis. I.e. if there are three links leaving that node the probability of selection of any one link based on the relative cost of the link. This probability selection is modified as the path search is developed as will be discussed below.

When an agent reaches the destination node it then returns to the originating node following the previously traversed links in reverse order as identified from its tabu list. As each link is traversed the agent leaves a volatile marker over the link, the value of that marker being inversely proportional to the total cost of the path. If more than one agent traverses a link then the values of their markers on that link are summed. The value of the marker may decay with time at a rate determined by the system operator. The technique ensures that the value of the marker on each link never becomes zero, as this will force random movement of the agents, according to the probability function below.

At this stage of the search further exploratory agents arriving at the node Ni (Fig. 2) are presented with a number of links L21, L22 .... some or all of which carry various levels M_i of the marker. Each of these links also carries a cost factor C_i. Selection of a link by which to leave the node and proceed to the next node is determined by the agent from a selection probability factor P_i as defined below:-

$$P_i = \frac{C_i^a M_i^b}{\sum_i C_i^a M_i^b}$$

where *a* represents the sensitivity of the agents to the cost of the link and *b* represents the sensitivity of the agents to the current level of the marker. This improves the results of the algorithm as this selection function acts as a greedy heuristic selecting the best solution. In this way, the lower cost paths will receive higher value markers and will this attract more and more new agents. The movement of the software agents from the source mode following their insertion can be summarised by the following algorithm:-

1.   Find possible links with sufficient bandwidth.

2.   Select one link based on probability factor.

3.   Add link to tabu list.

4.   Add link cost to total path cost.

5.   Move along link to next node.

6.   If next node = destination node, retrace path to originating node leaving marker. Stop on reaching originating node.

7.   If next node ≠ destination node go to step 1.

When the emergence of a preferred path is detected, this path is then allocated by sending an allocator agent from the originating node along the path to the destination node. This agent checks each path link to confirm that it can be allocated to the path and reserves a bandwidth partition on the path. If an allocation error is detected, the agent deallocates the link and returns to the originating node to update the connection. Having determined a path from the model, this path can then be set up on the network by the operator.

After a successful allocation of a path, the search process may be continued. At this point the available bandwidth on the allocated path will have been reduced by the allocation thus increasing its cost factor. This encourages the algorithm to seek alternative paths e.g. to identify a spare path that could be allocated in the event of a system fault, or to find a more cost effective path. In the latter event, the old path is deallocated and the new path becomes the allocated path

In a development of the technique, the effective cost of a link may be adjusted by the operator to take account of the usage level or channel occupation of the link and thus effect load balancing within the network. Typically we provide a significantly increasing cost penalty for load levels exceeding about 75% occupancy. A typical relationship between occupancy and cost of a link is illustrated in figure 3.

To demonstrate the technique described above, we have applied the process to the analysis of a model network as shown in figure 4. In this network, each link is allocated a cost factor as indicated by a respective integer adjacent that link. The process was used with various values of a and with a fixed value of b to determine a route between the nodes N1 and N9. In each case the optimum route found after a number of iterations was N1-N2-N6-N7-N9. The results of these tests are summarised in Table 1 below which shows the minimum and maximum number of iterations required to find the route for each set of values of a and b.

Table 1

a, b	$a=1, b=2$	$a=2, b=2$	$a=4, b=2$	$a=10, b=2$
Minimum	90	90	140	140
Maximum	333	220	204	200
Mean	212.45	170.82	160	154.4
Standard deviation	44.38	24.3	13.25	10.7

We have found from the above results that the optimum value of a is between 2 and 4, as this provides a reasonably rapid solution without the risk that the process will lock into the first found solution and will ignore possible better solutions.

In a further development to the technique, two sets or colonies of software exploratory agents are inserted, one at each of the end nodes of the path. This development further improves the technique by eliminating local optima due to the structure of the network.

In another embodiment, the behaviour of the exploratory agents may be developed using a genetic algorithm technique which encourages the agents to seek new paths after an initial path has been found and thus avoid the risk of locking on to a local optimum solution. This is achieved

by controlling the sensitivity of the agents to the marker and to the cost factor. when the initial set of agents is created, each agent is allocated cost and marker sensitivity parameters on a random basis. When an agent returns to the source node having discovered a path, its parameters are stored and linked to the cost of that path. This cost performs the role of a fitness function in the genetic algorithm. When new agents are created, the parameters of the recently returned agents are employed to generate an intermediate population of parameters. Also some random parameters are added to the population and the genetic operators such as mutation and crossover are carried out to provide a new set of parameters for the newly created agents. We have found that the introduction of the random parameters results in the exploration and exploitation of new paths thus increasing the probability of finding improved solutions.

In the above description, the route identifying algorithm has been described with particular reference to point to point routing. However it will be appreciated that it can be adapted for the determination of point to multipoint routing problems e.g. for applications such as distance learning where a single source node transmits information to and receives information from a number of destination nodes.

CLAIMS:

1.   A method of operating a telecommunications network whereby to determine traffic routing, said network comprising a plurality of nodes interconnected by communications links and each said link incorporating one or more communications channels, wherein the method includes, providing a map or model of the network, determining in the model those nodes between which a communications path or route is to be established across the network, creating a set or colony of exploratory agents at a said node, permitting the agents at said node to explore the model network via the links and nodes, recording the results of said exploration, and determining from said results a route favoured by a significant number of said agents, and setting up said route corresponding to that path on the telecommunications network.

2.   A method of operating a telecommunications network whereby to determine traffic routing, said network comprising a plurality of nodes interconnected by communications links and each said link incorporating one or more communications channels, wherein the method includes, providing a map or model of the network, determining in the model first and second nodes between which a communications path or route comprising a chain of links is to be established across the network, allocating to each said link a cost value representative of the operational cost of that link, creating a set or colony of exploratory agents at said first node, permitting the agents at said first node to explore the model network via the links and nodes between said first and second nodes, determining for each said agent the nodes and links visited by that agent whereby to inhibit more than one traversal of a said link by that agent, recording the results of said exploration by providing on each traversed path a marker whose value corresponds to the number of agents selecting that path, and determining from said results a route favoured by a significant number of said agents and setting up said route on the telecommunications network, wherein each said agent arriving via a link at a network node has means for selecting a link by which to leave that node on a probability based choice weighted with the respective costs of the respective links whereby to select preferentially a link of lower cost.

3.   A method as claimed in claim 2, wherein the total cost of each path traversed by the agents between the first and second nodes is

determined and a marker whose value is inversely proportional to that total cost is associated with that path.

4. A method as claimed in claim 3, wherein a further set of exploratory agents is created at said second node.

5. A method as claimed in claim 3 or 4, wherein said agents are created as successive generations, and wherein the response of said successive generations of agents to the cost of a link and the value of a marker on a path is developed via a genetic algorithm.

6. A method, of operating a telecommunications network substantially as described herein with reference to and as shown in the accompanying drawings.

7. Apparatus for determining traffic routing in a telecommunications network comprising a plurality of nodes interconnected by communications links, and each said link incorporating one or more communications channels, the apparatus including, means for providing a map or model of the network, means for determining in the model first and second nodes between which a communications path or route is to be established across the network, means for creating a set or colony of exploratory agents at said first node, said agents being adapted to explore the model network via the links and nodes between said first and second nodes, means for recording the results of said exploration, and means for determining from said results a route favoured by a significant number of said agents whereby to set up said route on the telecommunications network.

8. Apparatus for determining traffic routing in a telecommunications network substantially as described herein with reference to and as shown in the accompanying drawings.

Patents Act 1977 ⬤	_10_	A ⬤ ation number
Examiner's repor⬤ the Comptroller under Section 17 (The Search report)		GB 9519613.5

Relevant Technical Fields	Search Examiner MR A STRAYTON
·(i) UK Cl (Ed.N) H4K: KFM ·	
(ii) Int Cl (Ed.6) H04M, H04Q	Date of completion of Search 28 NOVEMBER 1995
Databases (see below) (i) UK Patent Office collections of GB, EP, WO and US patent specifications. (ii) ONLINE: WPI; INSPEC	Documents considered relevant following a search in respect of Claims :- ALL

Categories of documents

X: Document indicating lack of novelty or of inventive step.

Y: Document indicating lack of inventive step if combined with one or more other documents of the same category.

A: Document indicating technological background and/or state of the art.

P: Document published on or after the declared priority date but before the filing date of the present application.

E: Patent document published on or after, but with priority date earlier than, the filing date of the present application.

&:   Member of the same patent family: corresponding document.

Category	Identity of document and relevant passages	Relevant to claim(s)
A	GB 2253970 A (GPT)	
A	EP 0583135.A2 (ATT)	
A	US 5216591 (NEMIROVSKY)	
A	US 5067148 · (SARDANA)	

Databases:The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).

T8 - 25297   Page 1 of 1

This Page Blank (uspto)